

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau



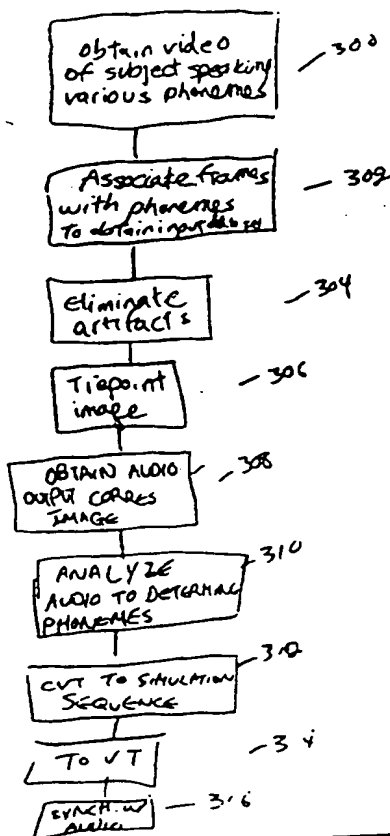
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>6</sup> : G06T 13/00, 7/20, 15/20, G10L 5/02, 5/04		A1	(11) International Publication Number: WO 96/17323
			(43) International Publication Date: 6 June 1996 (06.06.96)
(21) International Application Number: PCT/US95/15507		(74) Agent: HARRIS, Scott, C.; Fish & Richardson P.C., 601 Thirteenth Street, N.W., Washington, DC 20005 (US).	
(22) International Filing Date: 30 November 1995 (30.11.95)		(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TT, UA, UG, US, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG), ARIPO patent (KE, LS, MW, SD, SZ, UG).	
(30) Priority Data: 08/351,218 30 November 1994 (30.11.94) US			
(60) Parent Application or Grant (63) Related by Continuation US 08/351,218 (CIP) Filed on 30 November 1994 (30.11.94)			
(71) Applicant (for all designated States except US): CALIFORNIA INSTITUTE OF TECHNOLOGY [US/US]; 1201 East California Boulevard, Pasadena, CA 91125 (US).		Published With international search report.	
(72) Inventors; and (75) Inventors/Applicants (for US only): SCOTT, Kenneth, C. [US/US]; 2906 Fairmount Avenue, La Crescenta, CA 91214 (US). YEATES, Matthew, C. [US/US]; 2748 Montrose Avenue #2, Montrose, CA 91020 (US). KAGELS, David, S. [US/US]; 100 Hurlbut Street #13, Pasadena, CA 91105 (US). WATSON, Stephen, Hilary [US/US]; 254 South Madison #26, Pasadena, CA 91101 (US).			

(54) Title: METHOD AND APPARATUS FOR SYNTHESIZING REALISTIC ANIMATIONS OF A HUMAN SPEAKING USING A COMPUTER

(57) Abstract

A method and apparatus for synthesizing speech or facial movements to match selected speech sequences. A videotape of an arbitrary text sequence is obtained including a plurality of images of a user speaking various sequences. Video images corresponding to specific spoken phonemes are obtained (300). A video frame is digitized from that sequence which represents the extreme of mouth motion and shape (302). This is used to create a data base of images of different facial positions relative to spoken phonemes and diphthongs (306). An audio speech sequence is then used as the element to which a video sequence will be matched (308). The audio sequence is analyzed to determine spoken phoneme sequences and relative timings (310). The database is used to obtain images for each of these phonemes and these times, and morphing techniques are used to create transitions between the images (204). Different parts of the images can be processed in different ways to make a more realistic speech pattern.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

- 1 -

**METHOD AND APPARATUS FOR SYNTHESIZING REALISTIC  
ANIMATIONS OF A HUMAN SPEAKING USING A COMPUTER**

**Field of the Invention**

5       The present invention defines techniques allowing a computer to simulate an animated image of a human speaking. More specifically, the present invention uses special techniques to simulate human facial expressions associated with various speaking patterns.

10       **Background and Summary of the Invention**

Computer animation has been used to produce computer-generated pictures associated with various characteristics. Usually a computer animation is used to produce a moving animational system. As the users speak, 15 their mouths move, but the movement of the mouths of the speakers and their speech has not been synchronized. This does not bother the viewer, however, since it appears to be a cartoon; and is not intended to be accurate.

20       The inventors of the present invention recognized that usual computer animation does not provide a sufficiently accurate picture of a user speaking to allow it to be used as a facsimile of that user speaking. That is, under the current state of the art, the inventors of 25 the present invention recognized that a viewer of the computer animation would never be fooled into believing that the computer animation was real. They set about trying to find a way to solve this problem.

The inventors recognized, for the first time, that 30 morphing technology could be used to simulate moving facial characteristics. Morphing technology is well known in the art: it is used to simulate a continuous change from a first image of a first object into an image

- 2 -

of a second object. For example, it is easy to morph an apple into an orange. While one is looking at the apple, one sees its characteristics gradually change. It gradually assumes the shape of the orange, and also  
5 gradually assumes the color of the orange.

Morphing is well known, but a brief explanation of its operation will be given here anyway. Morphing involves transforming a first object, an "original object" into a second object, a "destination object".  
10 The computer takes the original object and the destination object, and maps various points thereof. These points define the shape and contour of both objects as well as the colors at the various points. Morphing can be carried out using a number of different  
15 techniques. For simplicity, we can assume that a small number of points, e.g., 16 points are used.

The morphing process is then calculated in advance: an interim point between the two objects is calculated, and then interim points between those objects  
20 are calculated. These interim points can be any points between the two objects. This provides a plurality of images, each image differing from the previous image by only a small amount, and each image incrementally closer to the destination image. By providing a number of  
25 images, over an amount of time, the difference between each two adjacent images is very small. The viewer sees the illusion of transformation from one image to another, and thus the user sees a continuously-varying image that changes gradually from the original image to the  
30 destination image. It appears as though the apple changes into the orange.

The inventors of the present invention were the first to realize that such morphing technology could be used to simulate an image of a human body part moving

- 3 -

between a first position and a second position for computer animation purposes.

One specific aspect of the present invention is the use of morphing technology to change facial image characteristics in a way to simulate the characteristics of speech. To do this, the inventors of the present invention developed a plurality of tools which changed human facial expression in accordance with speech to be spoken. The detailed aspects will be described herein.

It is therefore an object of the present invention to provide a system and method which changes facial expressions of a user's body part, preferably a user's face, in a way that associates those facial expressions with speech.

#### Brief Description of the Drawings

These and other aspects of the present invention will be described in detail with respect to the accompanying drawings, in which:

Figure 1 shows a block diagram of the hardware used according to the present invention;

Figure 2 shows a flowchart of operation of the first embodiment of the present invention;

Figure 3 shows a flowchart of operation of the second embodiment of the present invention;

Figure 4 shows a sample path editor without interpolation;

Figure 5 shows an example animation between keypoints;

Figure 6 shows a flowchart of operation of the grouping;

Figure 7 shows a plurality of points defining a boundary;

- 4 -

Figure 8 shows an additional point added to the boundary, and the effect that adding this additional point has on the boundary; and

Figure 9 shows a non-linear animation path between 5 keypoints.

#### Description of the Preferred Embodiments

The system includes multiple embodiments, which represent progression between the various stages of complexity of the invention. The first embodiment is  
10 described herein with reference to Figures 1 and 2.

The preferred mode of this invention contemplates forming an animation sequence of face and head and shoulders of a subject speaking. More generally, however, the present invention could be used to form an  
15 animation sequence of any action taken by the subject using the same concepts as described herein. While all description is given for speech and facial movements, of course, this teaching could easily be adapted to any movement.

#### 20 First embodiment

Figure 1 shows a basic block diagram of the hardware used according to the invention. The subject  
100 is located at a position where its image can be acquired by an image acquisition device 102, preferably a  
25 video camera with an associated image digitizer. The image digitizer can be an A66 and A25 device available from Abekas Video systems, Inc, Redwood City Ca. The Abekas Video Tools software allows the user to control the A66 from an SGI computer. The output of the image  
30 acquisition device is connected to a storage unit such as a video recorder or a dual port RAM or the like, and also to a computer 110, preferably a UNIX-based computer with an associated memory 112. The preferred mode uses a

- 5 -

Silicon Graphics (SGI) Indigo2 Extreme. Alternately, the hardware used can be more dedicated hardware, such as an image processing digital signal processor ("DSP"), or even by LSI circuitry.

5           The storage device 104 is also connected to a monitor 120, from which video images indicative of those taken by the image acquisition device 102 may be displayed. The computer 110 may also have a direct connection to the monitor 120.

10           For the first embodiment of the invention, the computer executes the flowchart of Figure 2. At step 200, an image of a subject with a first facial expression is acquired and stored. This image is preferably a view of at least the head of the subject, and preferably the  
15 head and shoulders of the subject. For purposes of this example, we will assume that the person is saying the word "rain". Step 200 then acquires an image of the person saying the "r" part of the word rain. At step 202 a second facial expression is acquired. In this example,  
20 the second facial expression is an expression of the user saying the sound "n" to end the word rain, and showing the distortions of their facial expressions caused by the speaking.

          At step 204, the computer interpolates  
25 intermediate images in an animation sequence which represent images which are produced from the first image and the second image. Preferably, this is carried out using commercially-available morphing software to morph between the images: from the first facial expression to  
30 the second facial expression. The morphing is carried out over a time period equivalent to that required for the word rain to be reproduced. It is also carried out in synchronism with the user saying the word "rain". This results in the user saying the word rain,



- 6 -

simultaneously with the morph between the first to the second facial expressions.

By morphing along the path between the various images, an image is obtained corresponding to the audio track.

Of course, this example given above omits two sounds/facial expressions between r & u.

Another aspect of the initial embodiment was to simulate a blink and/or a smile using morphing techniques similar to those discussed above. This used a first image of eyes open ("image 1"), a second image of half-way between eyes open and eyes closed ("image 2"), and a third image of eyes closed ("image 3"). The eye blink is then morphed by morphing image 1->2->3, holding it there, then 3->2->1. The smile can be morphed in a similar way.

This first embodiment required some trial and error, and also produced some distortion of facial features, since the user's features are not terribly natural in this state. This first embodiment, therefore, produced a rudimentary operation with sufficient realism to be usable, but having some problems therein.

#### Second Embodiment

The second embodiment uses more sophisticated tools to provide processing improvements in the facial expression. It provides geometrical perspective changes as part of the morph between images. The second embodiment also uses tools which allow better registration of the images to produce a more realistic final image.

The second embodiment breaks the speech to be simulated into units - specifically phonemes and/or diphthongs. Diphthongs are a type of phoneme. Images of

- 7 -

each unit are obtained, and the system morphs between the obtained images or keyframes.

Phonemes are the primary components of speech. There are many phonemes in the English language, the number of which varies depending on the way they are counted. For purposes of this application, we assume there are about 50 phonemes. In addition to phonemes, human speech also includes diphthongs. The inventors found that people characteristically change different parts of facial expressions -- face shape, mouth shape, head shape -- in different ways depending on the presence of the phonemes and diphthongs. For the second embodiment, the ways that the face changes were determined by trial and error.

The process of speech simulation according to the present embodiment carries out the flowchart of Figure 3. The process begins by acquiring images corresponding to a list of phonemes of the language. Ideally, this is obtained from a video of a person speaking over a certain time. Phonemes are identified within the speech, either by manual manipulation, or by the use of the ABEKAS (TM) video tools available from Abekas Video systems, Inc, Redwood City Ca. Each of the phonemes is associated with a frame that is determined to best fit the phoneme. Each of those frames is then captured, numbered and stored on disk, to form a first database which includes an entire set of phoneme images comprising the input data set, at step 302.

The speaker database is the fundamental representation of the figure to be animated in the output video. The database is a set of pictures of the subject's head/face, each picture having been digitized from the source video of the speaker. The synthesis process as used herein allows for complex combinations of database records to be used in the production of an

- 8 -

output picture, thereby increasing the possible output set to combinations of database elements.

The figure is represented in the speaker database as a set of digital picture of an actual person. Each  
5 picture is a record in the database. The various records in the database represent articulation of the face over the range of face shapes desired to be reproduced in the synthesized video sequence. For speech-related articulation, each record corresponds to the production  
10 of a phoneme. Other records may relate to other facial characteristics such as eyelid motion (open and closed), eyeball look direction (up, down, left, and rights), and emotion.

Figure representation is based on the visible  
15 speech model which relates to a set of speech related records of face shape in the speaker database to the production of a spoken phoneme. The input to the model is a sequence of spoken phonemes, the output is a sequence of database records or combination of records  
20 that reproduce the correct face shape during phoneme utterance.

The initial visible speech model as described herein expresses this relationship as one-to-one, i.e. each spoken phoneme is represented by one unique face  
25 shape in the database. The phonemic coding scheme uses 50 phonemes. The various phonemes which are preferably used according to this embodiment are as shown herein in table 1, but it should be understood that any other definitional organization of phonemes could alternately  
30 be used.

- 9 -

TABLE I

<u>symbol</u>	<u>Example words</u>
a	wad, dot, odd
b	bad
5 c	o in "or", au in "caught", aw in "awe"
d	add
e	angel, blade, way
f	farm
g	gap
10 h	hot, who
i	long e as in "eve", theme, bee
k	cab, keep
l	lad
m	man, imp
15 n	gnat, and
o	only, own
p	pad, apt
r	rap
s	cent, ask
20 t	tab
u	boot, ooze, you
v	vat
w	we, liquid
x	a in "pirate", o in "welcome"
25 y	yes, senior
z	zoo, goes
A	long i as in "ice", height, eye
C	chart, cello
D	the, mother
30 E	many, end, head
G	length, long, bank
I	i in "give", u in "busy", ai in "captain"
J	jam, gem
35 K	anxious, sexual
L	evil, able
M	chasm
N	shorten, basin
O	oil, boy
40 Q	quilt
R	honer, after, satyr
S	ocean, wish
T	thaw, bath
U	wood, could, put
45 W	out, towel, house
X	mixture, annex
Y	use, feud, new
Z	s in "usual", s in "vision"
@	cab, plaid
50 !	z in "nazi", zz in "pizza"
#	x in "auxiliary", x in "exist"

- 10 -

\*           wh in " what"  
 ^           u in "up", o in "son", oo in "blood"  
 +           oi in "abattoir", oi in "mademoiselle"

Certain substitutions in the last few phoneme characters  
 5 can avoid confusing the Unix operating system. We  
 suggest:

	aa	cab, plaid
	zz	z in "nazi", zz in "pizza"
	xx	x in "auxiliary", x in "exist"
10	ww	wh in "what"
	uu	u in "up", o in "son", oo in "blood"
	oi	oi in "abbattoir", oi in "mademoiselle"

At step 304, the input data set is processed to eliminate  
 artifacts. Various artifacts affect the realism of the  
 15 final simulated image. According to this embodiment, the  
 removed artifacts include lighting inconsistencies, small  
 amounts of subject motion as the subject is speaking,  
 variations in camera output over time, and the like.  
 These preprocessing operations use well-known image  
 20 processing functions including histogram equalization,  
 image registration and the like.

The next step in the process tiepoints the images  
 to one another. The tiepointing occurs among various  
 images of the input data set. Tiepointing is the process  
 25 of matching the specific features in one image with an  
 identical feature in another image -- even if that  
 feature is in different locations. For example, the  
 user's eyes, lips, teeth and hair may be tiepointed.  
 This embodiment requires the user to manually select the  
 30 points to be tiepointed in one image ("the reference  
 image"). The system automatically finds the same  
 tiepoints in all of the multiple images as described  
 herein. The input data set is completely set up once all  
 the images are set and tiepointed.

- 11 -

At step 308, the audio output, to which an image is to be synchronized, is obtained. This can simply be a user recording the audio sequence; it can be synthesized from the original sound/video sample which was used to form the image database. At step 310, that audio is analyzed to determine the phonemes which correspond thereto. One way which this can be done is by obtaining a written transcript of the audio track, and using a computer dictionary to determine the phonemes formed by the words in that written transcript. At step 312, the determined phonemes are converted into a simulation sequence.

Speech related facial motion in the present invention is based on interpolation of face shapes. Face shapes are stored as a set of control points for each picture in the database. The control points identify the location of facial features for each face shape. Of particular importance to speech are the facial features that vary in the production of speech. The main feature is the mouth, which includes the lips, teeth, tongue, jaw, and cheeks.

The speaker database must contain the set of face shapes over which the face need range to visually simulate speech. The major component of face shapes in the database are, as described above, pictures of the subject speaking a full set of phonemes. The visual appearance of speech is produced by displaying in order and at the appropriate rate a sequence of face shapes based on a phonemic translation of the desired speech. To smooth the motion of the figure, the interval between face shapes is filled with frames synthesized by morphing from the face shape at the beginning of the interval to that at the end.

For example, if the word to be spoken is "Poe", translated as /p/ occurring at time A and /o/ at time B,

- 12 -

in the range of time between A and B the mouth will linearly transition from the pursed lip shape the bilabial /p/ to the lip-rounded shape of the /o/. The /p/ picture is displayed at time A, the frames between A and B are synthesized from a linear combination of A and B by morphing, and the /o/ picture is displayed at time B.

The initial visible speech shape model must fully express the face shape of each phoneme. The simulation occurs from linearly interpolating between phonemes. At maximum acoustic expression of the phoneme, the relevant face shape in the speaker database fully controls the face shape in the synthesized output video. Test sequences shows that full visual expression of all phonemes has an unnatural appearance. Visually this results in unnaturally fast, jerky and extreme mouth motion.

The visible speech model is modified to base extent of visual expression on the location of sound production in the vocal tract. Generally, a sequence of phonemes is established that controls the shape of the face. The controlling phonemes are produced mainly by the lips and teeth. Phonemes produced behind the teeth in the mouth cavity affect the shape of the face without controlling it. Phonemes produced behind the vellum have no control or affect on face shape. The affect on face shape is accomplished by establishing keyframes that are a linear combination of face shapes, the major percentage from the controlling phoneme and the minority percentage from the affecting phoneme. This will be discussed further herein.

Sample video sequences have shown that certain phonemes do not have an associated face shape (i.e. face shape is irrelevant to the production of the sound) while others may have influence on face shape without

- 13 -

controlling it. Also, visual expression of certain phonemes is sensitive to the context of the preceding and succeeding phonemes.

After the complete processing, a full resolution output data set is generated at step 314, and saved to memory. These output frames are then transferred to the Abekas digital/analog converter and then to video tape at step 314. At step 316, this is synchronized with the audio track.

10 The following describes the steps of the flowchart of Figure 3 in further detail. Step 300 is the initial preparatory step of obtaining a video of a subject speaking various phonemes. This video can be obtained from the subject conducting a dedicated session to speak  
15 the various phonemes, or by obtaining a video tape showing the subject speaking, e.g. a speech or news broadcast. This source must be converted into a set of phoneme images which comprise the input data set. Ideally, there should be controlled studio-like setting  
20 with proper lighting control, consistent video white balance, and a smooth and relatively featureless background. If any of these characteristics are not available, then the database portion should choose images which are as nearly identical in subject position and  
25 orientation as possible. Discrepancy between head position and different images causes a less stable final product. However, head movements within plus or minus 5 degrees in any of the three axes of motion still allow a quite acceptable product.

30 Once an appropriate portion of the video tape is chosen, the frames must be associated with the phonemes of step 302. This is done using the Abekas system. Abekas allows the user to jog through the simulation frame by frame. The user can then determine which of the  
35 frames best matches the phoneme. As described above, the



- 14 -

best match is usually obtained when the facial features reach their maximum movement. The Abekas software is used to manually scan through the segments to locate images which represent the various phonemes. In this  
5 embodiment, the user must manually analyze each phoneme image to ensure that it is the most correct image among the several frames which correspond to the phoneme. That most correct frame is usually the one at which the mouth, teeth, and tongue are at the most extreme positions  
10 relative to adjacent images. These most extreme positions enables the best end point for the morphing.

The Abekas system copies appropriate frames into the computer memory. A table is formed in memory, correlating each phoneme to a frame to which it  
15 corresponds. When a frame is selected as being representative of a phoneme, that frame number is noted in the table to correspond to the phoneme. A database is accordingly established in memory between the frame number and the phoneme.

20 Once all phonemes are entered into the database, we have one frame corresponding to each phoneme at Step 302. Step 304 then processes these frames to minimize the artifacts so that the final images will be consistent. Then the morph between the images provides a  
25 more realistic final animation. This artifact correction includes color correction and registration.

Color correction is a process of adjusting relative values of the images in the data set to compensate for variations in camera output values. A  
30 single image is chosen as a reference image. The remaining data set is manipulated so that its range of values best matches that reference image. This is done by obtaining histograms of various characteristics of the image values, including their color saturation and the

- 15 -

like. The color histograms of all the other images are equalized to the reference.

Registration is the process of positionally aligning one image to another. The registration  
5 according to the present invention first selects two points which will remain fixed in one image, e.g., the user's eyes and nose. The rest of the data set is then registered to that image by rotating and translating the rest of the images comprising the data set. Other  
10 choices of the fixed points include the inside corners of the eyes. This effectively ties each of the images to the same positional system. After completion of this pre-processing, the image is tiepointed at step 306.

Tiepoints are positions - x, y locations in an  
15 image - which correspond to a feature in the image. These points are correlated with the same points in another image and "tensioned" with respect to that other image. The tensioning affects the amount of movement of tiepoints that can occur between any images -- e.g. a  
20 tension of 0 allows unlimited movement while a tension of 1 holds the tiepoints firmly to one another. Suitable examples of tiepoint locations include corners of the eyes, pupils, selective points around the irises, eyebrows, lips, teeth, hair and the like. Any point  
25 which is distinct from its nearby region can be used as a tiepoint.

Tiepointing thus allows portions of the images to be associated with one another in a controllable way. Tiepointing is normally carried out on the entire outline  
30 of the subject, working around the shoulders, neck and head of the subject. Usually, these head-outline tiepoints will be fairly strongly tied to one another to prevent random bobbing motion of the head as the user speaks. This set of boundary tiepoints must be fairly  
35 dense to ensure a good outcome.

- 16 -

The inventors found that the ultimate tiepointing density for these features is every 20 pixels or so.

Additional tiepointing is necessary for those facial features which move during speech. This includes, for example, the eye shapes, mouth shapes, etc. One tiepoint every 5-10 pixels has been found optimum for this.

Teeth, tongue, and eye-balls pose the additional problem of being occluded during portions of the animation. These features are tiepointed with regard to the use of groups to allow them to selectively be made to appear and disappear in a natural fashion.

Finally, we need to tiepoint additional features such as the nose, cheeks, chin, neck and the like to ensure a reasonably uniform set of tiepoints. The synthesizing algorithm, as described herein, triangulates among these points, making it important to tiepoint as many features as possible to assure the best triangulation. This also makes the final product more realistic by ensuring that the features change in a relatively smooth and localized manner. Tiepointing can cause unusual artifacts due to triangulation if the triangles that are used are too large. Too dense a triangulation, in contrast, makes the synthesis creation process slower. 200-300 tiepoints per image has been found by the inventors to be optimum. The preferred technique of triangulation operates as follows.

The triangulation method used in this system is the Cline-Renka Generalized Delaunay Triangulation (GDT). The GDT is a generalization of the standard Delaunay Triangulation (SDT) which can deal with non-convex regions, holes, and edge constraints. The SDT problem can be stated as follows: Given a (finite) set S of points (nodes) in the plane, determine a set T of triangles such that:

- 17 -

- 1) The vertices of the triangles are nodes,
- 2) No triangle contains a node other than its vertices,
- 3) The interiors of the triangles are pairwise disjoint,
- 4) The union of the triangles is the convex hull of S,
- 5 and
- 5) The interior of the circumcircle of each triangle contains no nodes.

Property (1) is an obvious requirement.

Properties (2) and (3) prevent triangles from  
10 overlapping. Property (4) ensures that the entire region  
in question is covered by the triangulation.

Property (5) can be shown to be equivalent to the  
optimality condition of maximizing the minimum angle in  
the triangulation over the set of all possible  
15 triangulations. An SDT is equivalent to a Dirichlet  
tessellation and  
to a Voronoi diagram.

The preferred algorithm for the solution of the SDT is as  
follows:

- 20 1) Create an auxiliary triangle A so that S is entirely  
contained in A. Add A to T.

- 2) For each point p in S:

- 2a) Find the set T' of triangles whose  
circumcircle contains p.

- 25 2b) Determine the union of the triangles T',  
called the "insertion polygon", I.

- 2c) Find the outer (boundary) edges of I.

- 2d) Create new triangles T'' by connecting p to  
the vertices of I.

- 30 2e) Delete T' from T, and add T'' to T.

- 3) Remove all triangles which share a vertex with A.

Other algorithms for creating the SDT without  
using an auxiliary triangle exist, but the increased  
computational complexity was not worth the gain in this

- 18 -

case. This other algorithms are usually used when some outside factor prevents the capability of setting up the auxiliary triangle.

Now, having computed an SDT over our set of points  
5 S, we want to ensure that certain required edges and boundaries are included in the triangulation. However, if these edges are added to the triangulation by some means, then property (5) will be lost, and the optimality of the triangulation is usually no longer true. Thus,  
10 some sort of modified circumcircle property is needed. Note also that if we specify interior boundaries, then property (4) will no longer be true. The following definitions and modifications are made to provide for these boundaries and required edges:

15 Definition 1: Let  $B = \{B_i, i \geq 1\}$ , where the  $B_i$ 's are simple, closed polygonal curves in the plane, pairwise disjoint. The line segments composing each  $B_i$  are called "Boundary Edges".

Definition 2: Let  $\Omega = \text{closure}(\text{interior}(B_1) \cap \text{interior}(B_2) \dots)$ .

Definition 3: Let E be a set of "Required Interior Edges". Required interior edges are line segments which connect pairs of nodes. No other nodes lie on the line segment, and the line segments are interior to  $\Omega$ .

25 Now, E union B constitutes the set R of "Required Edges".

We can now modify the circumcircle test, property (5), by weakening it as follows:

5') For any triangle t in T, if some node is  
30 contained in the interior of the circumcircle of t, then that every interior point of the triangle t is separated from that node by a required edge.

This means that a triangle can pass the circumcircle test even if some node is inside its  
35 circumcircle, but only if the node in question lies "on

- 19 -

the other side" of some required edge, i.e., the triangle is on one side of a required edge, and the required edge is one side of the triangle, and the node in question lies on the other side of the required edge.

5 We now have the following Generalized Delaunay Triangulation (GDT) problem to solve:

Given a (finite) set  $S$  of points (nodes) in the plane, a set of polygonal boundary curves  $B$  which define  $\Omega$ , and a set of required edges  $E$ , determine a set  $T'$  of  
10 triangles such that:

- 1) The vertices of the triangles are nodes
- 2) No triangle contains a node other than its vertices
- 3) The interiors of the triangles are pairwise disjoint
- 4') The union of the triangles is  $\Omega$
- 15 5') If any node is contained in the interior of the circumcircle of a triangle, then every interior point of the triangle is separated from the node by an element of  $R = E \cup B$
- 6) Each element of  $R$  is an edge of at least one triangle.

20 And the Cline-Renka solution to the GDT is as follows:

- 1) Determine the SDT,  $T$
- 2) For each edge  $e$  in  $R$ , call `add_edge(e, T, R)`
- 3) Delete all triangles with interiors exterior to

25  $\Omega$

Procedure `add_edge`(edge  $e$ , triangulation  $T$ ,  
required\_edge\_list  $R$ )

- 1) Find the triangles in  $T$  whose interior intersect edge  
 $e$

30 1a) If no such triangles exist, stop;  
else, remove all such triangles from  $T$

2) Let:

$\Omega_e$  = set of all triangles of step 1  
 $B_e$  = boundary edges of  $\Omega_e$

- 20 -

$R_e$  = any required edges in  $\Omega_e$   
nodeset = set of nodes in  $\Omega_e$  other than  
endpoints of  $e$   
 $a, b$  = endpoints of  $e$

5 3) Retriangulate the "left" side of the required edge  $e$   
retriangulate(nodeset,  $a, b, T, R_e$  union  $B_e$ )

4) Retriangulate the "right" side of the required edge  $e$   
retriangulate(nodeset,  $b, a, T, R_e$  union  $B_e$ )

5) Replace  $R$  with  $R$  union  $e$

10 and the real work occurs in:

Procedure retriangulate(list nodeset, point  $p_1$ ,  
point  $p_2$ ,  
triangulation  $T$ , required\_edge\_list  $R$ )

1) Find all the nodes strictly left of the line  $(p_1, p_2)$ ,  
15 which are not separated from the midpoint of the line  
 $(p_1, p_2)$  by some other required edge. Denote this set  $X$ .

2) Find the node  $x$  in  $X$  that maximizes the angle  $p_1-x-p_2$ .

3) Add the triangle  $(p_1, x, p_2)$  to  $T$

4) Delete  $x$  from nodeset

20 5) If the line from  $p_1$  to  $x$  is not in  $R$ ,  
retriangulate(nodeset,  $p_1, x, T, R$ )

5) If the line from  $x$  to  $p_2$  is not in  $R$ ,  
retriangulate(nodeset,  $x, p_2, T, R$ )

Conceptually, this operates as follows.

- 21 -

- 1) Compute an SDT
- 2) For each edge that has to be added, create an insertion polygon for that edge, along with some assorted arrays of nodes, etc.
- 5 3) Find all the points in the insertion polygon that are strictly left of the edge, and retriangulate so that the modified circumcircle test will be met. Then, using the two new sides of the triangle just created, recursively retriangulate the remaining points.
- 10 4) Do the same thing for all the points that are strictly right of the edge.
- 5) Clean up by removing any exterior triangles. This completes the triangulation.

After the tiepoints are selected for one

15 particular image, a matching algorithm applies these same tiepoints to the other images. This embodiment carries out the matching by investigating the corresponding locations in the other images. The pixel areas around these corresponding locations are then correlated against

20 the pixels forming the tiepoint in the originally tiepointed image. The best correlation between areas is taken as the corresponding tiepoint.

The matching algorithm is at least 90% effective in selecting locations of the features in the images.

25 Ideally, therefore, each of the images in the database should be investigated to ensure that the matching algorithm has properly placed the tiepoints. This is preferably done manually.

Preferably, a table stores information about the

30 tiepoints, including a tiepoint identifier, which can be a number, for example, and the x,y coordinates of that tiepoint. The tiepointed image is displayed as the image, overlaid with the tiepoints from the table. The operator then investigates the image to manually

35 determine if the tiepoints are placed in their proper



- 22 -

locations. If not, a tiepoint editor can be used to change the x,y coordinates for each tiepoint. In its simplest embodiment, the tiepoint editor is simply an editor which calls up the table and changes the x,y  
5 information associated with one of the tiepoints therein.

Once the tiepoints have been established for each image, the database has been established. This database can now be used to produce a simulation or animation, sequence. See step 312 at Figure 3. This is done using  
10 a tool that the inventors have called "the animator". The animator uses the various databases to produce an animation sequence of the user speaking using the tiepoint images, and the phoneme images.

The animation is defined by keyframes at specific  
15 points in time. Each keyframe is a point in time which is described fully. All times between keyframes are not described fully; they are simulated images that are simulated from parts of the keypoint, or known, images.

The keyframes can be defined from one image or  
20 from multiple superimposed images. The keyframes also include a plurality of tiepoints in the image.

The animation follows a path between keyframes. That path is interpolated between the known data which exists at the keyframe. The same path may be used for  
25 both the images and the tiepoints or alternately separate paths may be used.

For example, some sounds, such as "p" affect the shape of the face. This causes the positions of the tiepoints forming the face shape to change. The "ha"  
30 sound causes the face look to change, and affects the throat shape. The "a" sound comes from the middle of the mouth. These sounds and their associated shapes show that different sounds affect different face/mouth parts etc, differently. The paths of the tiepoints and the  
35 images therefore differ for these elements.

- 23 -

It is often useful to have several images at each keyframe. The portion of each image to be used at a particular keyframe can be defined. Different combinations of images and tiepoints at given times vary  
5 the realism and the look of the final animation. The final image is formed of a linear weighted combination of the images.

The tiepoint tensioning and pathing allows different portions of the animation to be separately  
10 controlled. The eyes can be moved independently of the mouth, for example. The animation begins by defining a path. For example, the best path may be a path for the head.

An example path editor screen is shown in Figure  
15 4. This path editor screen, shows a timeline of the animation. The initial timeline is shown in Figure 4 with each of the keyframe times being shown for each of a plurality of phonemes. A simple animation between the images a, b, E, etc. is shown in Figure 5.

20 As described above, multiple images can be used at a keyframe. This is done by selecting multiple images for a keyframe and the relative proportions of the images at the keyframe. If there are several images, the lines connecting the images represent the various phonemes  
25 which will be combined.

One additional tiepointing feature is the boundary/grouping operation. The boundary/grouping operation of the present invention begins with a completely tiepointed database. The images and tiepoints  
30 are grouped by defining boundaries for each group.

For this first embodiment, each of the defined groups include all of the tiepoints in the database, obviating the need to specify which tiepoints are in each group. The boundary of the group allows the morphing and

- 24 -

animating software to automatically exclude portions of the image which are outside of that boundary.

According to the present invention, each of the boundaries includes a group name and a group level.

- 5 Lower- numbered levels take precedence over higher numbered levels in making the final image. We will give herein an example of the group "head," level 0." Figure 6 shows the flowchart of operation.

- 10 As explained above, we begin by selecting the group name and group level at step 600. For this example we have chosen "head" as the name and 0 for the group number.

- At step 602, we define the boundary of the head group by selecting the tiepoints in the tiepointed image which correspond to the boundary of the head. This boundary should be a closed surface which encloses various points. Additional boundary curves are preferably defined at step 604; and the group is therefore defined between an inner boundary curve and an outer boundary curve.

- The first boundary is preferably the outside of the head, with the second boundary excluding the eyes and mouth. The multiple boundaries define a group wherein the morphing algorithm excludes all areas outside of the outer boundary and they also excludes all areas inside the inner boundary. The boundary is stored in memory as a series of points defining the area therein. The computer determines a connection between the points, preferably a plurality of separated line segments, to define the boundary. These points, and hence this boundary, may also be edited at step 606 to add additional points, for example. Figure 7 for example, shows a selected boundary comprising a plurality of tiepoints. To change this boundary, an additional point

- 25 -

has been added in Figure 8. Points may also be deleted in this same way.

Next, other groups may be added, such as an eye group or the like.

5       The purpose of these groups is to allow various parts of the image to be animated independently of the other parts. In order to do this, the boundaries of these groups must match throughout the animations. In addition, certain points of the subject must be held  
10 still during the animation.

For example, holding the shoulders still during the animation makes a more realistic product.

Tensions for the groups are defined beginning at step 608. These tensions are numerical values between 0  
15 and 1 as described above. The value of the tensions determine how closely a tiepoint is held to either a reference image or a reference path. Each group is defined to have a separate tension and a separate reference path. However, the boundaries between groups  
20 must be held to a common path or else gaps would appear between the groups during the morphing process. Therefore, tension values may be assigned to the outer edges of the boundary to maintain that boundary line.

Normally we set the lower shoulder tiepoint values  
25 and groups close to 1.0 in order to keep the shoulders from moving in an unnatural fashion during talking. We then proceed up the shoulders towards the neck where we set tension values progressively lower. Most of the other tiepoints are set to 0.0 or another low value,  
30 since movement of these other tiepoints makes the animation more lifelike.

The groups and tension files in the tiepointer can also be used in the animator. An example is set forth herein. We use the example discussed above - creating  
35 two groups: one for the head with a boundary surrounding

- 26 -

the entire head but excluding the eyes, and a boundary around the eyes. Another group includes the eyes - with the outer boundary being at the same boundary as within the head group. Preferably we set the head group at or  
5 close to level 0, and the eyes group at or close to level 1. This means that the head morphing will be overlaid onto the eye morphing. More generally, the layers are overlaid in order from back to front, higher to lower. The head does not cover the eye group since the inside  
10 boundary of the head leaves a hole through which the eyes can be seen.

In order to create a lifelike image, we must hold those tiepoints to the path to which the head is being morphed. Even if the eye group has its own path which is  
15 completely independent of the phonemes which are being synthesized, the boundaries will still match up. This is important, since the eyes can and do move entirely independent of the mouth. Of course, we can set the boundary points to 1.0 in both groups, in which case the  
20 head and the eyes will be held to completely separate paths.

### Third Embodiment

Certain aspects are further improved in the third embodiment.

25 First, the third embodiment further improves the tiepointing by improving the automated matching operations between images once tiepoints are chosen.

Each of the matching operations attempts to determine the location of the best match between images  
30 by comparing a small region around each original tiepoint in the original image with similarly-shaped regions in the target images. A correlation between the two regions is computed and the center of the specific region with the best correlation being used as the matched tiepoint

- 27 -

location. A number of preferred techniques effect the matching operation.

A first technique computes integer pixel matching. This is the fastest but least-accurate technique. This  
5 technique computes the correlation between regions of different images by using only integer movements within the correlation area. This technique is therefore accurate to only a single pixel.

The other correlation operations are labelled as  
10 models 0-4 are various implementations of the so-called Gruen Subroutine

This is a discussion of the different correlation options supported under the Gruen subroutine. Five  
options exist, called modes 0 to 4. Each option will be  
15 discussed below.

#### Mode 0:

This mode corresponds the closest with what most people consider to be correlation. In this case the template is kept on integer pixel boundaries and is  
20 matched to each possible pixel location in the search area. That integer location which matches the best is selected for the next step. In order to return a sub-pixel location, the correlation values of points on a column and a row through the best match point are fitted  
25 to a quadratic and the peak of the quadratic is selected and returned in both sample and line. Since the correlation is never actually performed at the sub-pixel location, the returned values are only estimates. The returned correlation value is still that determined at  
30 the best integer location. If the peak correlation occurs at pixel  $i$  and has an amplitude  $f(i)$ , then the location of the interpolated peak is computed from

- 28 -

$$x = i - \frac{1}{2} \frac{f(i-1) - f(i+1)}{2f(i) - f(i-1) - f(i+1)} \quad (1)$$

and similarly for y.

Execution time per tiepoint depends upon the template size and the search area size. For a typical area it requires a normalized time of 1.0.

5           Mode zero is the mode of choice when one knows:

1.           The rotation and scale differences between the template and the search area are minimal.
2.           That the location of the best correlation could be anywhere in the search area.

10          3.           That accuracy is never required to exceed 1/10 of a pixel.

4.           That time is of the essence.

Mode 1:

15           This mode makes use of simulated annealing to arrive at the best correlation location. This is the slowest of all options computationally, and the most experimental. It was provided as a means of last resort when other options have been exhausted and a correlation is still required.

20           The method uses guessing the six polynomial coefficients to an affine transformation which maps the template onto the search area. Any mapping is acceptable provided it remains within the search area. Each guess adds to the last location six values obtained from a  
25   random number generator constrained to remain within a certain range or "temperature". Temperature, here, is a metaphor borrowed from the source of the title for this method, "Simulated annealing." Annealing is "heating and then cooling to ...". In this case the temperature  
30   refers to the range of numbers used. As the temperature of the mathematical system is lowered, it settles upon a solution to the six parameter equation. If the system is cooled too fast, it may yield an incorrect answer.

- 29 -

Gradually the temperature is reduced so that guesses remain more localized. The heart of the technique is to compute, at each step, the Boltzmann probability of transition from the previous correlation  
5 to the current one. If the current correlation is higher than the last then we adopt the new affine position. If it is lower we compute the probability of that transition and compare it with chance. This is analogous to accepting the transition if the coin comes up, and  
10 rejecting the transition if the coin comes down. The essence of annealing is that it gives us a way of escaping from local false minima in the solution space. Thus, this can be considered as a non-deterministic method because the next move is not constrained entirely  
15 by the last move. Unlike Mode zero, Mode 1 does not systematically search the solution space for all combinations of mappings. It starts at an initial estimate and bounces about trying all sorts of combinations of affine mapping while remembering the best  
20 location visited. Repeatedly it is forced to revisit the best correlation location. Gradually the angle of guesses is reduced until it freezes near the best minimum. If the number of iterations is kept small it will freeze at the wrong location. If the number of  
25 iterations is kept large the best correlation location will always be found but at the expense of time. Several thousand iterations should be used.

This mode requires more arguments than any of the others. It requires the Gruen the input mapping  
30 polynomial coefficients Line\_coef\_limits and Samp\_coef\_limits; the temperature ranges on each of the coefficients Line\_temp and Samp\_temp; and the iteration limit Limits.



- 30 -

Execution time per tiepoint depends upon the template size and the number of iterations. For a typical area it requires normalized time of 14.0.

Mode one is the mode of choice when one knows:

- 5 1. That these is an unknown amount of rotation, scale, or distortion between the template and the search area.
2. That the location of the best correlation could be anywhere in the search area.
- 10 3. That if the images are distorted this distortion is to be compensated.
4. That accuracy is never required to exceed  $1/30$  of a pixel. Actually the user can control this precision.
5. That time is unimportant in exchange for a
- 15 tiepoint.

Mode 2:

This mode uses the simplex downhill search strategy. A simplex is a tetrahedron with one corner greater than the dimension of the problem or surface it resides on. In this case the surface is one where the correlation value is a function of six dimensions, those of the six affine mapping coefficients. These coefficients map the template to the search area. We seek the six coefficients which map the template to the search area, and for which  $(1 - \text{the correlation})$  is a minimum. The simplex stands on the surface. There are four rules describing permitted changes in shape for the simplex as it seeks to move along the surface towards a minimum. Eventually it will find the bottom of the correlation surface and will compress itself down to the desired precision. This is a deterministic method because the next move depends entirely upon the last.

Deterministic schemes have the drawback that if they start in the wrong minimum they have no means of

- 31 -

escape. Therefore, the initial estimate for the affine mapping polynomial must be within the correct minimum for Mode two to function correctly. The user can control the starting location for the search but this only sets two  
5 of six affine coefficients. In most cases the driving program provides initial estimates for a unity mapping. This is adequate if the initial tiepoint is within the correlation distance. If data were strongly distorted, however, it might not suffice. Mode two does not search  
10 the entire search area. It begins at one location and events guide it from there.

This mode requires an input estimate of the mapping polynomial, arguments Line\_coef, and Samp\_coef in Gruen.

15 Execution time per tiepoint depends only upon the template size. For a typical area it requires a normalized time of 1.0.

Mode two is the mode of choice when one knows:

1. That if there is rotation, scale, or distortion  
20 between the template and the search area that initial mapping polynomial coefficients are available to begin a search within the correct minimum.
2. That the initial tiepoint location is within a few pixels of the true one.
- 25 3. That as much accuracy is desired as possible.
4. That if the images are distorted this distortion is to be compensated.
5. That time is important but subordinate to accuracy.

30 Mode 3:

This is hybrid mode. In this case Mode zero is first used to determine the tiepoint location. This location is then passed on to Mode two along with a unity mapping transformation. Since Modes zero and two are  
35 comparable in execution time, this is a good combination.

- 32 -

It provides broad search and great precision in the result.

Execution time per tiepoint depends upon the template size and the search area size. For a typical  
5 area it requires a normalized time of 1.8.

Mode three is the mode of choice when one knows:

1. That the location of the best correlation could be anywhere in the search area.
2. That the amount of rotation, scale, and distortion  
10 between the template and the search area is slight.
3. That accuracy is essential.
4. That time is important.

Mode 4:

This is also a hybrid mode. In this case Mode one  
15 is used first to determine the mapping coefficients. These coefficients are then passed to Mode two which refines the solution. This combination provides great precision along with the minimum of a priori knowledge on the part of the user.

20 Execution time per tiepoint depends upon the template size and the number of iterations. For a typical area it requires a normalized time of 16.0.

Mode 4 is the mode of choice when one knows:

1. That the location of the best correlation could be  
25 anywhere in the search area.
2. That the amount of rotation, scale, and distortion between the template and the search area is substantial or unknown.
3. That accuracy is essential.
- 30 4. That time is unimportant.

The Object Function

Each of the five correlation modes discussed above is really a means for determining the location in an image where some quantity, which we call correlation  
35 value, is a maximum. This quantity is computed in the

- 33 -

same fashion for all modes and is itself mode independent. Since it is a scalar to be optimized, it really is an objective function. Gruen uses a least squares objective function called the coefficient of determination. It measures the quality of a least squares linear fit made between the intensity values of the template and the corresponding intensities in the search area as determined by the affine mapping polynomial. This objective function value lies between 0.0 (no correlation at all) to 1.0 (correlation or anti correlation), and is returned as argument Quality is subroutine Gruen. The correlation quality is computed from

$$r^2 = \frac{(\sum xy - \frac{\sum x \sum y}{mn})^2}{(\sum x^2 - \frac{(\sum x)^2}{mn})(\sum y^2 - \frac{(\sum y)^2}{mn})} \quad (2)$$

where x and y are the intensity values in the template and the search area respectively. Note that because the measure is a least squares determination, correlation quality is indifferent to intensity differences between the template and the search area which are of the nature of scale, offset, or complement. Anti-correlations are just as valid as correlations since both imply non randomness.

All modes except Mode 0 permit the template to suffer a distortion compared with the search area. The nature of the distortion is anything that a first order polynomial or affine transformation can do. There are six coefficients involved, three for sample and three for line. The modes one through four are concerned with determining these six coefficients. By varying the coefficients, one can simulate changes in offset, scale, rotation, skew, transpose, or flip. Since we are really

- 34 -

interested in the tiepoint location, we only want the offset term in the sample and the line equations, however, we need to compute all the terms in order to extract the desired terms.

- 5       The third embodiment also allows adding small rotational movements to the images associated with the synthesized speech to thus create a more realistic change in perspective while simulating usual head movement. This provides a simulation of depth information and  
10 allows tiepoints to be moved along the Z-axis e.g. to include depth information added therein.

- Depth information may be added to either a single tiepoint or to a group of tiepoints by selecting the tiepoint or group and selecting the amount of depth  
15 information to be added therein.

- The rotation information is added by simulating the look of an image rotation. Assuming the head is the shape being simulated, we need to model the three dimensional shape of the head. This model tells us the  
20 two dimensional look of the head shape when looking from the front of the head, and from various angles.

- Now, we added some random rotational movements to the head to make it look more natural. Most speakers move and slightly rotate their heads when speaking.  
25 Random movements in the z direction therefore help the realism. These z movements change the shape of the head according to the model described above.

- If there is depth information added to the tiepoint data, then rotations can be added to the final  
30 animation. According to this embodiment, translations can also be added to the final information whether or not depth information is present. According to this embodiment, some rotation is defined on the x, y, and z planes. Usually the same "curves" indicative of the  
35 rotation are used for every group. The rotation should

- 35 -

be kept below approximately + or - 5 degrees. Any further rotation results in artifacts.

As described above, the images are only fully defined at the times referred to herein as keyframes.

5 Between the keyframes, the images are interpolated along the path. The first and second embodiments linearly interpolate between the keyframes, using morphing techniques. As described above, multiple images may be defined at any one keyframe, and the outputs would  
10 correspond to summations of these images.

Figure 9 shows the advanced animating features available according to this embodiment. These advanced animating features enable a non-linear interpolation between keyframes: essentially the "morph" becomes non-  
15 linear. One example of such a non-linear morph is shown in Figure 9. Non-linear interpolation may also be used for motion. More specifically, the spline-based paths are used to control translation and rotation of groups. The linear paths are used to control the tiepoint and  
20 images transitions. Non-linear interpolation is defined in terms of the tiepoints and images in that extra keyframes which can be inserted between the keyframe associated with phonemes to provide for non-linear paths that are piece-wise linear.

25 The first and second embodiments operate according to path 800, a linear morph between keyframe A and keyframe B. At the 50% point between keyframe A and keyframe B, point 802, the image is composed of half of the keyframe A and half of the keyframe B. This produces  
30 a smooth transition between images, and has many advantages. However, one problem with this system is the signature that it leaves on the final product. If one were to investigate the frames, one would find a linear transition between keyframe images. Such linear  
35 transitions would be very unlikely to occur in nature.

- 36 -

The non-linear transition according to the third embodiment therefore enables a non-linear morph between images. An example non-linear morph is shown as path 804. This path can follow any function whatsoever or can  
5 be entirely random. At the 50% point between keyframes A and B, point 806, the image is much closer to B than it is to A. If one investigates the images between keyframes, one finds a non-linear pattern which can be, for example, random.

10 According to this embodiment, additional post-processing is possible in the output frames. This can be used to add in background images, color corrections, blur the edges or add noise or the like in order to improve the realism of the final product.

15 One technique which the inventors have found to be very useful in improving the final product is to add gaussian noise to the image. This decreases the quality of the final image. However, unexpectedly, it also makes the image look more realistic by hiding some of the image  
20 parts that detract from its realism.

The gaussian noise which is used as a pseudo random gaussian noise produced, for example, by a UNIX computer. A window size for the gaussian noises is selected, and the noise width is set via a convolution  
25 kernel. The size and pixel spreads out based on the noise width, with the noise strength representing the amplitude of the noise.

A "composite" function blends the foreground image with a background image based on an alpha value of the  
30 foreground image. This is used, for example, if the background that was used in the original production is not sufficiently suited for the animation. Some backgrounds would show, for example, the movements necessary to register the images. The composite function

- 37 -

can be used to remove the existing background, and substitute a new background therefor.

The composite function assigns alpha ( $\alpha$ ) values to more than one image. The alpha channel defines the amount of transparency of an image. An alpha image with a value of 0 is not transparent. Therefore, by setting the head and shoulders to an alpha value of 0, the head and shoulders will always show over a background image. The background itself is set to an alpha value of maximum here,  $2^8=255$ . This alpha value renders the background completely transparent. Therefore, everything behind this image can be seen through the transparent image.

The composite function then carries out a linear addition of pixels. A non-transparent pixel always shows through a transparent pixel. This allows the background image to be added outside the set boundaries.

Another advanced feature in the third embodiment is the ability to control the attack and decay of the face shapes. Face shapes change during speech. The face shape is defined by the outer parameters of the head; thus the image defined by the outer parameters of the head changes during the morph between keypoints. According to this aspect of the invention, the face shape changes according to an attack and decay function. Each face shape changes toward its destination face shape with an attack function. It changes away from its destination shape following a decay function.

The third embodiment also uses a defocus function, which applies a convolution to the input image in order to produce a blurring or defocusing affect.

The edge blur function applies a convolution similar to that of defocus but only to the edges of the foreground image. The defocus function applies a convolution to the input image to produce a blurring or defocusing effect.



- 38 -

Based on the initial visible speech model, a full speaker database was produced and a set of animations was synthesized to demonstrate the resulting level of realism. An analysis of the results and examination of the database revealed that diphthongs are not adequately represented by a single face shape and that the database contains redundancies in face shape.

The production of a diphthong acoustically is a glide between two sounds. The start and end sounds are approximately that of two vowels, as a plot of F1 versus F2 formants clearly shows. Thus, visually, the shape of the face must also be represented as a glide between two face shapes. The visible speech model was extended to include representation of diphthongs as a glide between two face shapes, represented by the records in the speaker database corresponding to the production of the relevant two vowels. Sample video sequences were produced to test this hypothesis; the result was more realistic expression of the face shape to accompany the sound of a diphthong.

The speaker database of face shapes contained obvious redundancies. Two approaches to reducing the redundancies have been considered. First, eliminate redundancies based on characteristics of productions such as voiced/unvoiced pairs and location in vocal tract. Second, categorize the face shapes and eliminate commonality. Sample video sequences were synthesized based on substitution of voice/unvoiced pairs with no appreciable visual difference. Reduction based on categorized face shapes have not yet been tested.

Other tools can also be used to improve the realism of the final image.

Although only a few embodiments have been described in detail above, those having ordinary skill in the art will certainly understand that many modifications

- 39 -

are possible in the preferred embodiment without departing from the teachings thereof.

All such modifications are intended to be encompassed within the following claims.

- 40 -

What is claimed is:

1. A method of producing a computer-based animation of a subject, comprising:

obtaining a first image of a first position  
5 of the subject which represents a starting point of a first part of an animation;

obtaining a second image of the subject at a second position of said first part of said animation;

forming a database of images including at  
10 least said first and second images;

relating some aspect of each of the images in the database to the same aspect in each of the other images in the database; and

establishing an animation sequence between  
15 said first image and said second image, while maintaining a specified relationship between said aspects to produce an animation of the subject moving from said first position to said second position while maintaining said specified relationship between said aspects.

20 2. A method as in claim 1 wherein said images of said subject are images of a user's face and head at various points in speech, said speech is formed of units, said first and second images and other images in said database corresponding to said different units of said  
25 speech.

3. A method as in claim 2, wherein said animation is an animation of the subject speaking, and said images are images of facial expressions corresponding to said units of speech.

30 4. A system for producing a computer-based animation sequence of a subject, comprising:

- 41 -

an image storage device, storing a database of images of a subject, including at least a first image of a first position of the subject which represents a starting point of a first part of an animation, a second  
5 image of the subject at an ending point of said first part of said animation and a plurality of aspects in said images which relate to one another;

a computer, connected to receive information indicative of said images and said aspects from said  
10 image storage device, and interpolating intermediate images which also have a predetermined relationship with said aspect, in an animation sequence which represent images which are produced from said first image and said second image; and

15 a display unit, connected to display said animation sequence including, in order, said first image, said intermediate images, and said second image.

5. A system as in claim 4 wherein said computer morphs between said first and second images via said  
20 intermediate images, while maintaining said aspects of said images

6. A method of producing a computer-based animation of a subject speaking, comprising:

determining a set of units of speech;  
25 preparing a database of images, each said image corresponding to one of said units of speech;  
establishing some aspect of each image of said database which relates to each other image in the database;  
30 obtaining a sequence of speech to which said animation is to be synchronized;  
analyzing said sequence of speech to determine said units of speech therein;

- 42 -

determining keyframe images which correspond to said units determined by said analyzing; and

using said keyframe images to produce said animation sequence by defining relationships among the  
5 aspects in a way that maintains at least one of said aspects in a predetermined relationship with another of said aspects.

7. A method as in claim 6 wherein said aspects are tiepoints, and said using includes interpolating  
10 between said keyframe images to maintain the desired relationship between said tiepoints based on an amount of tieing therebetween said tieing being variable between a first value that requires tiepoint positions to be precisely at the same positions relative to one another,  
15 a second value that allows the tiepoint positions to move freely relative to one another, and at least one value in between said first and second values that holds said tiepoints to one another by a specified amount less than said precisely but more than said freely.

20 8. A method as in claim 7, wherein said units are at least one of phonemes or diphthongs, and said database of images includes images corresponding to said phonemes and/or diphthongs.

9. A method as in claim 6 wherein said units are  
25 units of speech, and said preparing includes:  
obtaining a sample of the subject speaking;  
investigating frames of the sample to identify the units of speech therein;  
determining a frame of the sample which best  
30 represents a particular unit of speech; and  
storing an image in the database representative of said frame of the sample which best represents the

- 43 -

particular unit correlated with an indication of the particular unit.

10. A method as in claim 6 wherein said images are images of a user's head and face.

5           11. A method as in claim 9, further comprising, prior to said storing, forming an image indicative of said part of said sample, and preprocessing said image to make it more consistent with other images in the database.

10           12. A method as in claim 11, wherein said preprocessing includes changing an amount of lighting effect in an image to equalize a lightness of the image to other images in the database.

15           13. A method as in claim 11, wherein said aspects include positions of portions of one image and said defining includes registering in the database relative to positions of the same portions of other images in the database.

20           14. A method as in claim 6, wherein said aspects are tiepoints, comprising determining tiepoints in each image which are associated with similar tiepoints in other images in the database, and simulating motion between said images, while controlling a position of said tiepoints in said images relative to one another.

25           15. A method as in claim 9, wherein said aspects are tiepoints, said establishing comprises identifying tiepoints in the images, wherein said images in the database each include tiepoints associated therewith, and maintaining a desired relationship between said tiepoints

- 44 -

based on an amount of tying therebetween said tying amount being variable between a first value that requires said tiepoint positions to be precisely maintained at the same positions relative to one another, a second value  
5 that allows the tiepoint positions to move freely relative to one another, and at least one value in between said first and second values that holds said tiepoints to one another by a specified amount less than said precisely but more than said freely.

10        16. A method as in claim 9, wherein said aspects are tiepoints and establishing comprises identifying a plurality of tiepoints in a reference image, and matching said reference image to all other images in the database to identify tiepoints in said images.

15        17. A method as in claim 16, wherein said matching comprises identifying locations in said images which correspond to areas of said reference image where said tiepoints are located.

20        18. A method as in claim 17, wherein said matching step comprises correlating said areas of said reference image to areas of said other images, by determining an area of said image which maximally matches to an area of said tiepoint of said reference image.

25        19. A method as in claim 6, wherein said analyzing comprises determining units in said sequence of speech.

20. A method as in claim 6, wherein said using comprises interpolating between face shapes.

- 45 -

21. A method as in claim 20, wherein said interpolating comprises interpolating between said keyframe images.

22. A method as in claim 6, wherein said  
5 preparing comprises preparing a table including a plurality of frame numbers associated with unit numbers.

23. A method as in claim 6, comprising  
determining two points in two images in the database  
which will remain positionally constant, and aligning a  
10 first image to a second image using said two points.

24. A method as in claim 23, wherein each said keypoint includes a plurality of images.

25. A method as in claim 24, wherein said using  
comprises producing an animation sequence by morphing  
15 between said keyframes.

26. A method as in claim 25, further comprising  
defining a plurality of groups in each of a plurality of  
images, and said using comprises morphing each of said  
groups separately.

20 27. A method as in claim 26, wherein each said group includes a group name and a group number, wherein said groups are animated separately.

28. A method as in claim 6, wherein said tiepoints include tensions associated therewith, said  
25 tensions controlling an amount by which said tiepoints are permitted to move relative to one another.



- 46 -

29. A method as in claim 6, further comprising means for allowing depth changes.

30. A method as in claim 7, wherein said interpolating is linear.

5 31. A method as in claim 7, wherein said interpolating is non-linear.

32. A method as in claim 6 wherein each keyframe is defined by two keyframe images, said two images include a first image with a first weighting, and a  
10 second image with a second weighting, said relative relationships defined by said weighting which define an amount of transparency relative to one another, such that an image with a first weighting is transparent relative to an image with a second weighting.

15 33. A method of producing a computer-based animation sequence of a subject carrying out an action, comprising:

obtaining a database of images, where each of said images corresponds to a unit part of said action;  
20 determining some aspect of each image in the database and relating said aspect in each image to the same aspect in another image in the database;  
obtaining information indicating said action;  
analyzing said information to determine which  
25 of said units exist therein;  
determining keyframe images which correspond to said units; and  
using said keyframe images to produce said animation sequence while maintaining a specified  
30 relationship between said aspects of images in the animation sequence.

- 47 -

34. A method as in claim 33, wherein said action is speech, said units are at least one of phonemes or diphthongs, and said information is a unit sequence of frames of images showing speech to be simulated.

5 35. A method as in claim 34, wherein said aspects are tiepoints, and further comprising maintaining a desired relationship between said tiepoints based on an amount of tying therebetween, said tying amount being variable between a first value that requires said  
10 tiepoint positions to be precisely maintained at the same positions relative to one another, a second value that allows the tiepoint positions to move relative to one another, and at least one value in between said first and second values that holds said tiepoints to one another by  
15 a specified amount less than said precisely but more than said freely.

36. A method of producing a computer-based animation of a subject speaking, comprising:

determining a set of units of speech, said  
20 set being one from which substantially all face shapes of the subject during speaking can be reconstructed;  
obtaining a recorded sequence of the subject speaking;

first analyzing said recorded sequence to  
25 determine frames within said recorded sequence which best correspond to said units of speech;

preparing a database of images using said frames within said recorded sequence which best correspond to said units of speech, said database  
30 including an identifier of said unit, and an identifier of said image, so that said database includes images corresponding to said units;

- 48 -

determining some aspect of each image in the database and relating said aspect in each image to the same aspect in another image in the database;

obtaining a sequence of speech to which said  
5 animation is to be synchronized;

second analyzing said sequence of speech to determine said units of speech therein;

determining keyframe images from said database of images, which correspond to said units of  
10 speech determined in said second analyzing; and

using said keyframe images to produce said animation sequence while maintaining specified relationship between said aspects.

37. A method as in claim 36, wherein said aspects  
15 are tiepoints, and said processor maintains a desired relationship between said tiepoints based on an amount of tying therebetween said tying being variable between a first value that requires said tiepoint positions to be precisely maintained at the same positions relative to  
20 one another, a second value that allows the tiepoint positions to move relative to one another, and at least one value in between said first and second values that holds said tiepoints to one another by a specific amount.

38. A method as in claim 36, wherein said using  
25 comprises non-linearly interpolating between said keyframe images.

39. A method of producing a computer-based animation of a subject taking an action, comprising:  
determining a set of units of said action,  
30 said set being one from which substantially all movements within the action can be reconstructed;

- 49 -

obtaining a recorded sequence of the subject performing many actions;

first analyzing said recorded sequence to determine frames within said recorded sequence which best  
5 correspond to said units of said action;

preparing a database of images using said frames within said recorded sequence which best correspond to said units of said action, said database including an identifier of said unit, and an identifier  
10 of said image, so that said database includes images corresponding to said units;

determining some aspect of each image in the database and relating said aspect in each image to the same aspect in another image in the database;

15 obtaining a sequence of said actions to which said animation is to be synchronized;

second analyzing said sequence of actions to determine said units of said action therein;

determining keyframe images from said  
20 database of images, which correspond to said units of said action determined in said second analyzing; and

using said keyframe images to produce said animation sequence.

40. A method as in claim 39, wherein said action  
25 is speaking, said aspects are at least one of diphthongs or phonemes, said sequence is a sequence of the subject speaking.

41. A method as in claim 39, wherein said aspects  
30 are tiepoints and further comprising maintaining a desired relationship between said tiepoints based on an amount of tying therebetween, said tying being variable between a first value that requires said tiepoint positions to be precisely maintained at the same

- 50 -

positions relative to one another, a second value that allows the tiepoint positions to move relative to one another, and at least one value in between said first and second values that holds said tiepoints to one another by  
5 a specified amount.

42. A method as in claim 39, wherein said using comprises non-linearly interpolating between said keyframe images.

43. An apparatus for producing a computer-based  
10 animation of a subject speaking, comprising:  
a database of images, each said image corresponding to one of a plurality of units of speech;  
means for obtaining a sequence of speech to which said animation is to be synchronized, said sequence  
15 of speech including said units of speech therein;  
a processor, operating to determine keyframe images which correspond to said units of speech present in said sequence of speech, and to use said keyframe images to produce said animation sequence by:  
20 a) obtaining a sample of the subject speaking;  
b) investigating the sample to identify the units therein;  
c) determining a part of the sample which best represents a particular unit; and  
25 d) storing an image in the database representative of said part of the sample which best represents the particular unit, and an indication of the particular unit; and  
an image preprocessor which forms an image  
30 indicative of said part of said sample, and preprocesses said image to make it more consistent with other images in the database.

- 51 -

44. An apparatus as in claim 43, wherein said processor interpolates between said keyframe images.

45. An apparatus as in claim 43, wherein said processor morphs between said keyframe images.

5 46. An apparatus as in claim 43, wherein said units are phonemes, and wherein said database of images includes an image corresponding to each phoneme of said set.

10 47. An apparatus as in claim 46 wherein said units also include diphthongs, and said database of images includes images corresponding to said diphthongs.

48. An apparatus as in claim 43 wherein said images are images of a user's head and face.

15 49. An apparatus as in claim 43, wherein said preprocessor changes an amount of lighting effect in an image to equalize a lightness of the image to other images in the database.

20 50. An apparatus as in claim 43, wherein said preprocessor registers positions of portions of one image to positions of other images in the database.

51. An apparatus for producing a computer-based animation of a subject speaking, comprising:  
a database of images, each said image corresponding to one of a plurality of units of speech;  
25 means for obtaining a sequence of speech to which said animation is to be synchronized, said sequence of speech including said units of speech therein;

- 52 -

a processor, operating to determine keyframe images which correspond to said units of speech present in said sequence of speech, and to use said keyframe images to produce said animation sequence, wherein said  
5 processor determines tiepoints in each image which are associated with similar tiepoints in other images, and simulating motion between said images, while controlling a position of said tiepoints in said images relative to one another.

10 52. An apparatus for producing a computer-based animation of a subject speaking, comprising:  
a database of images, each said image corresponding to one of a plurality of units of speech;  
means for obtaining a sequence of speech to  
15 which said animation is to be synchronized, said sequence of speech including said units of speech therein;  
a processor, operating to determine keyframe images which correspond to said units of speech present in said sequence of speech, and to use said keyframe  
20 images to produce said animation sequence; and  
means for identifying tiepoints in the images, wherein said images in the database each include tiepoints associated therewith.

25 53. An apparatus as in claim 52, wherein said tiepoint identifying means further comprises means for identifying a plurality of tiepoints in a reference image, and means for matching said reference image to all other images in the database to identify tiepoints in said all other images.

30 54. An apparatus as in claim 53, wherein said matching means comprises means for identifying locations

- 53 -

in said images which correspond to areas of said reference image where said tiepoints are located.

55. An apparatus as in claim 53, wherein said matching means comprises means for correlating said areas  
5 of said reference image to areas of said other images, by determining an area of said image which maximally matches to an area of said tiepoint of said reference image.

56. An apparatus as in claim 43, wherein said  
10 processor comprises means for producing an animation sequence by morphing between said keyframes.

57. An apparatus as in claim 56, wherein said processor further comprises means for defining a plurality of groups in each of a plurality of images, and  
15 morphs each of said groups separately.

58. An apparatus as in claim 57, wherein each said group includes a group name and a group number, wherein said groups are animated separately.

59. An apparatus as in claim 53, wherein said  
20 tiepoints include tensions associated therewith, said tensions controlling an amount by which said tiepoints are permitted to move relative to one another.

60. An apparatus as in claim 44, wherein said interpolating is linear.

25 61. An apparatus as in claim 44, wherein said interpolating is non-linear.



- 54 -

62. An apparatus for producing a computer-based animation of a subject speaking, comprising:

a database of images, each said image corresponding to one of a plurality of units of speech;

5 means for obtaining a sequence of speech to which said animation is to be synchronized, said sequence of speech including said units of speech therein;

a processor, operating to determine keyframe images which correspond to said units of speech present  
10 in said sequence of speech, and to use said keyframe images to produce said animation sequence;

wherein each said keyframe is defined by two images, said two images include a first image with a first weighting, and a second image with a second  
15 weighting, said weighting amounts defining an amount of transparency relative to one another, such that an image with a first weighting is transparent relative to an image with a second weighting.

63. An apparatus of producing a computer-based  
20 animation sequence of a subject carrying out an action, comprising:

a database of images, where each of said images corresponds to a unit part of said action, said database also including a relationship between an aspect  
25 of a first image in the database and a second image in the database;

means for analyzing an action to determine which of said units exist therein;

a processor which operates to determine  
30 keyframe images, from among said database of images, which correspond to said units; and

an animator, using said keyframe images to produce said animation sequence.

- 55 -

64. An apparatus as in claim 63, wherein said action is speech, said units are at least one of phonemes or diphthongs, and said indication of said action is a unit of speech to be simulated.

5           65. An apparatus as in claim 63, wherein said aspects are tiepoints, and said processor maintains a desired relationship between said tiepoints based on an amount of tying there between said tying variable  
10           between a first value that requires said tiepoint positions to be precisely maintained at the same  
            positions relative to one another, a second value that allows the tiepoint positions to move relative to one  
            another, and at least one value in between said first and  
15           second values that holds said tiepoints to one another by a specific amount

            66. An apparatus for producing a computer-based animation of a subject performing an action, comprising:  
            an image processing element, operating to  
            analyze the subject performing many aspects of the action  
20           to determine frames within said recorded sequence which best correspond to said units of said action and to  
            determine some aspect of the frames which are related to each other in a way to relate the frame to one another;  
            a database of images, formed from said frames  
25           within said recorded sequence which best correspond to said units of said action, said database including an  
            identifier of said unit an identifier of said image, and an identifier of the aspect so that said database  
            includes images corresponding to said units;  
30           means for obtaining a sequence of said actions to which said animation is to be synchronized,  
            and analyzing said sequence of actions to determine said

- 56 -

units of said action therein and to determine identifiers of images corresponding to said sequence, and

a processor, obtaining keyframe images from said database of images using said identifiers, which  
5 keyframe images correspond to said units of said action and using said keyframe images to produce said animation sequence by animating between said keyframe images while maintaining said desired relationship between said aspects.

10 67. An apparatus as in claim 66, wherein said action is speaking, said units are at least one of diphthongs or phonemes, said sequence is a sequence of the subject speaking.

15 68. An apparatus as in claim 66, wherein said aspects are tiepoints, and said processor maintains a desired relationship between said tiepoints based on an amount of tying therebetween, said tying variable between a first value that requires said tiepoint positions to be precisely maintained at the same  
20 positions relative to one another, a second value that allows the tiepoint positions to move relative to one another, and at least one value in between said first and second values that holds said tiepoints to one another by a specified amount.

25 69. An apparatus as in claim 66, wherein said processor non-linearly interpolates between said keyframe images.

70. A method as in claim 1, wherein said aspects are tiepoints.

- 57 -

71. An apparatus as in claim 70, wherein said aspects are tiepoints, and said processor maintains a desired relationship between said tiepoints based on an amount of tying therebetween, said tying being variable  
5 between a first value that requires said tiepoint positions to be precisely maintained at the same positions relative to one another, a second value that allows the tiepoint positions to move relative to one another, and at least one value in between said first and  
10 second values that holds said tiepoints to one another by a specified amount.




FIG 1




FIG 2

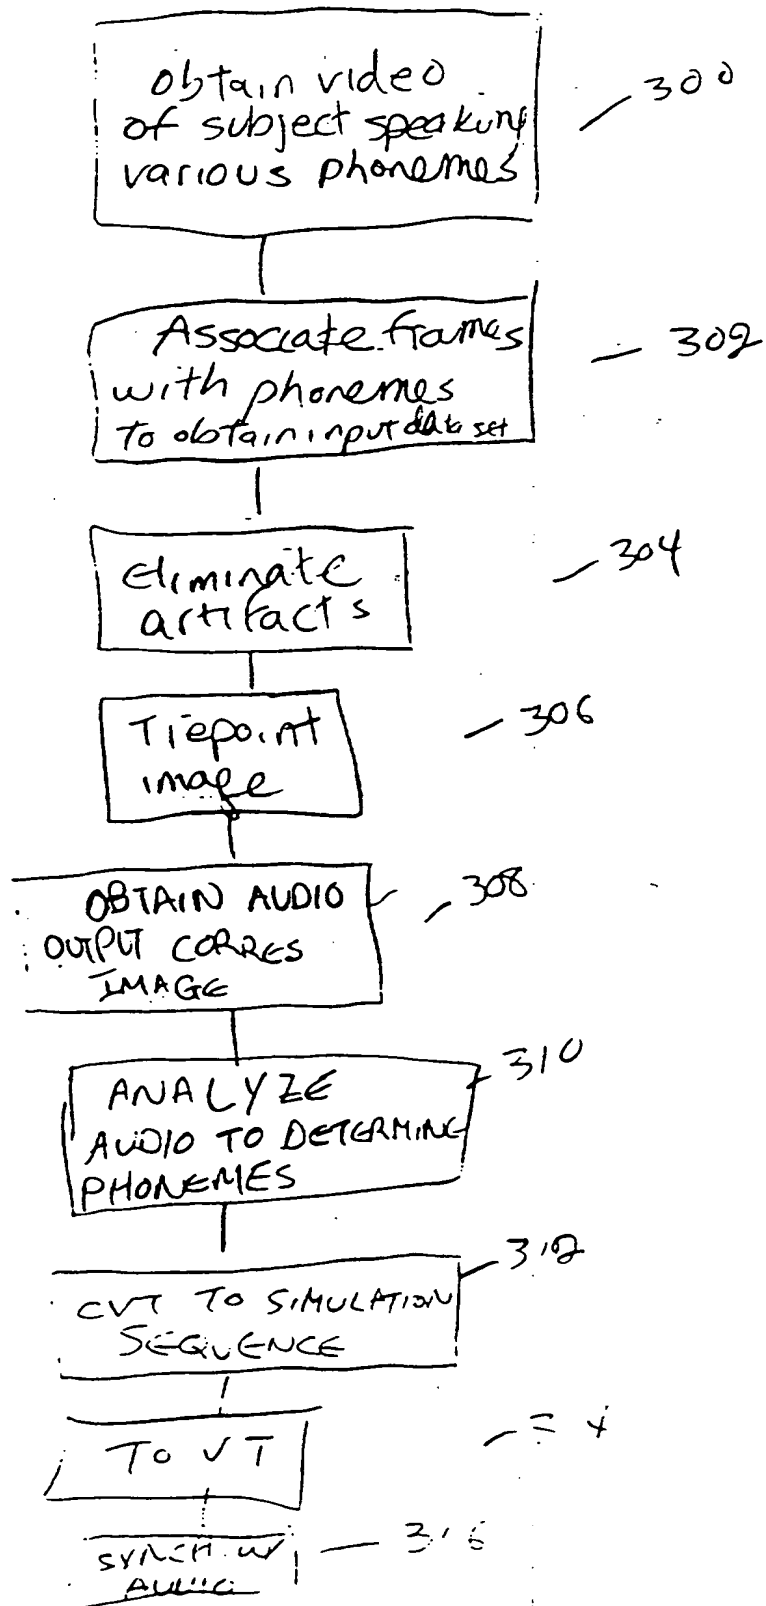


FIG 3




FIG 4





FIG 5



5/7






FIG 7




FIG 8




FIG 9

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US95/15507

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06T 13/00, 7/20, 15/20; G10L 5/02, 5/04

US CL : 395/2.84, 2.87, 129, 141, 152, 154

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/2.84, 2.87, 129, 141, 152, 154

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

IEEE Database;  
DIALOG

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	Computer Animation, 1994 Proceedings, September 1994, "LEE ET AL, "Image Morphing Using Deformable Surfaces", pages 31-39. see pages 31, column 2, lines 1-15 and paragraph 2; page 38, figure 6; page 31, column 2, paragraph 3, line 3.	1, 4 ----- 2-3, 5, 7-10, 15-20, 22-23, 25-30, 32-33, 35-71
X -- Y	Image Processing Neural Networks, 1994 International Symposium on Speech, April 1994, LAVAGETTO ET AL, "Lipreadable Frame Animation Driven by Speech Parameters", pages 626-629. see abstract, lines 10-12; page 628, paragraphs 1, 3	6, 24 ----- 2-3, 5, 7-23, 25-71

☒ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

•	Special categories of cited documents:	•T	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
•A	document defining the general state of the art which is not considered to be part of particular relevance		
•E	earlier document published on or after the international filing date	•X	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
•L	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	•Y	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
•O	document referring to an oral disclosure, use, exhibition or other means		
•P	document published prior to the international filing date but later than the priority date claimed	•Z	document member of the same patent family

Date of the actual completion of the international search

07 FEBRUARY 1996

Date of mailing of the international search report

04 MAR 1996

Name and mailing address of the ISA/US  
Commissioner of Patents and Trademarks  
Box PCT  
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

STEPHEN HONG

Telephone No. (703) 308-5465

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US95/15507

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	IEEE Transactions on Pattern and Machine Intelligence, vol. 15, no. 6, June 1993, TERZOPOULOS ET AL, "Analysis and Synthesis of Facial Image Sequences Using Physical and Anatomical Models", pages 569-579. see page 570, column 2, last paragraph and paragraph 2, line 4; page 571; page 572, figure 3; page 573, figure 4; page 574, column 2, line 5; page 576, paragraphs 3-4 and column 2, paragraph 2; page 577, paragraph 2, line 2 and figure 7.	11-15, 17-23, 25-71
A	US, A, 4,884,972 (GASPER) 05 December 1989, see entire document.	1-71
A	US, A, 5,111,409 (GASPER ET AL) 05 May 1992, see entire document.	1-71
A	Motion of Non-Rigid and Articulated Objects, 1994 Workshop, May 1994, ESSA ET AL, "Tracking Facial Motion", pages 36-42.	1-71
A	Image processing and Neural Networks, 1994 International Symposium on Speech, April 1994, BROOKE ET AL, "Computer Graphics Animation of Talking Faces Based on Stochastic Models", pages 73-76.	1-71
A	Acoustic Speech & Signal Processing, ICASSP '93, BOTHE ET AL, "Visual Speech and Coarticulation Effects", pages v634-v637.	1-71